

AD _____

Award Number: DAMD17-00-1-0244

TITLE: Improved Breast Cancer Research Through Automated Matching of Patients to Clinical Trials

PRINCIPAL INVESTIGATOR: Lawrence O. Hall, Ph.D.
Dmitry B. Goldgof, Ph.D.
Jeffrey Krischer, Ph.D.

CONTRACTING ORGANIZATION: University of South Florida
Tampa, Florida 33620-7900

REPORT DATE: August 2002

TYPE OF REPORT: Annual

PREPARED FOR: U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

DISTRIBUTION STATEMENT: Approved for Public Release;
Distribution Unlimited

The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision unless so designated by other documentation.

20021127 086

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	
	August 2002	Annual (3 Jul 01 - 2 Jul 02)	
4. TITLE AND SUBTITLE Improved Breast Cancer Research Through Automated Matching of Patients to Clinical Trials			5. FUNDING NUMBERS DAMD17-00-1-0244
6. AUTHOR(S) Lawrence O. Hall, Ph.D. Dmitry B. Goldgof, Ph.D. Jeffrey Krischer, Ph.D.			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of South Florida Tampa, Florida 33620-7900 E-Mail: hall@csee.usf.edu			8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Medical Research and Materiel Command Fort Detrick, Maryland 21702-5012			10. SPONSORING / MONITORING AGENCY REPORT NUMBER
11. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited			12b. DISTRIBUTION CODE
13. Abstract (Maximum 200 Words) <i>(abstract should contain no proprietary or confidential information)</i> An enhanced Web based prototype intelligent agent/expert system for matching breast cancer patients to clinical trials has been built. It allows for cost preferences to be entered. Therefore, the system user can choose to rule patients out of trials as quickly as possible without regard to the cost of tests necessary to do this. They can choose have questions appear so that the patient is ruled out of the trial with the minimal set of costs (tests) or can choose some combination of approaches. The system has been tested with 12 protocols and designed for maximal responsiveness and scalability as new protocols are added. The files of 178 former patients have been used to test the accuracy of the system. Additionally, the files of 57 current patients have been tested for eligibility. Patients for each of the protocols were correctly found eligible for one or more trials. We have also developed a prototype system to quickly add new clinical trials. This has been successfully used by novices to enter new trials.			
14. SUBJECT TERMS cancer control, clinical trials, expert system, breast cancer			15. NUMBER OF PAGES 19
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT Unlimited

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

Contents

1	Cover	1
2	SF 298	2
3	Introduction	4
4	Body	4
5	Conclusions	5
5.1	So What	6
6	Papers	8

3 Introduction

Increasing the enrollment of patients in clinical trials is important to making progress towards finding more effective treatments for breast cancer. Accrual is complicated by a large number potential studies and the cost and complexity of determining whether a patient meets the necessary eligibility criteria. Under this proposal, we are developing a Web based expert system which can determine the patients eligibility for clinical trials. The expert system is designed to take into account the cost of tests which are required to meet inclusion criteria and acquire information in the most cost-effective way possible.

Additionally, it is important to be able to easily add and remove clinical trials to the system. Trials are continually becoming available, going on suspension or being closed to accrual. Towards this end, we have developed a companion Web based system that enables anyone to simply enter the information required to describe the eligibility/ineligibility criteria for a clinical trial. A newly entered trial/protocol can then be directly included in the Clinical trial assignment expert system with no expert intervention.

4 Body

In this second-year, we have refined the original prototype to produce version 1.3.2. We have tested it with data from 187 retrospective patients and we have extensively tested its ability to order questions associated with tests to save dollar costs on 30 patients. We have further tested with 57 current patients and are continuing to test new patients as data becomes available. Table 1 summarizes our results on the current patients. Patients are only evaluated for trials that are currently enrolling patients. The trial status can change when a trial is put on suspension, closed, brought off suspension, or initiated. It can be seen that the system finds all matches that correspond to trials that patients have been enrolled in with one exception in which there is some missing data. Also, the 57 patients have been found eligible for 37 trials on which they were not enrolled. Clearly, a set of patients who are eligible for clinical trials are not being enrolled for some reason (s) (there are 28 in this class). On the day this is written, we have 74 current patients that have been run through the system.

We have verified that the system correctly finds protocols for which patients are eligible. We are investigating the cases where the system finds patients eligible for a protocol but they do not go on the protocol. These patients fall into two classes: the class of patients put on a different protocol and the class of patients not put on any protocol. There are now 12 protocols available in the system. Some of these are closed. At the present time, all breast cancer protocols at the Moffitt Cancer Center which are accruing at least two patients a month are available through our system.

It is important to be able to add new trials/protocols in a time efficient manner. It is also important that the process be such that it is straightforward for a physician or nurse or medical worker to enter the information from the eligibility/ineligibility criteria. Initially, it was taking us approximately one week of the time of a computer science expert to enter a new protocol. This year we have developed a prototype system which enables a user to enter a new trial/protocol in about an hour. We have tested it with novice users [1, 2] and found that they learn to use the system quite quickly. It is our conjecture that anyone with a modicum of medical knowledge

and access to the eligibility/ineligibility criteria or inclusion/exclusion criteria can enter a new protocol.

Key Research Accomplishments:

- We have enhanced our prototype system to very stable version 1.3.2. We have added cost functionality and tested this successfully.
- Utilizing retrospective patient data and current patient data, it has been found that patients are eligible for multiple protocols/trials. Further, with current patient data we find patients eligible for trials and not put on any trial.
- Extensive testing of cost functionality has been done. It has been shown that the average cost of determining eligibility may be significantly reduced (by over 60%) when the cost functionality portion of the extra system is utilized.
- An automated protocol acquisition tool has been developed, version 1.7. It has been tested with novice users and found to be quite usable. It is now how we add new protocols. A new protocol takes about 1 hour to enter.

Reportable Outcomes: We have had two papers [2, 3], which are attached, accepted to the 2002 IEEE International Conference on Systems, Man, and Cybernetics. We are in the process of revising a journal submission that got reasonably positive reviews. A web prototype of the clinical trial assignment system is available at <http://morden.csee.usf.edu/moffit> with password available from the principal investigator.

5 Conclusions

We have developed a scalable prototype which currently can determine eligibility for twelve breast cancer clinical trials. The system has been tested using retrospective data from 187 patients who are assigned to some clinical trial. Its accuracy has been verified. The system correctly finds cases in which a patient is eligible for multiple clinical trials. This will enable a physician to make the best choice from available trials. The system is able to utilize monetary cost in requesting tests to rule in/rule out a patient from the set of available clinical trials. The default ordering of questions allows the system user to rapidly determine the eligibility or ineligibility of a patient for any subset of the available clinical trials entered into the system. We have been able to show a significant average cost saving (over 60%) by using the cost feature to order questions. Of course, there is no guarantee that a clinician would order tests as suggested by the question ordering of our system. However, the potential for cost savings is significant.

The system is Web based and password protected. It provides rapid response when a person enters answers to one or more questions on a page of system selected questions. It can be used from any computer on the World Wide Web. Hence, community physicians will be able to determine the potential eligibility (they may not wish to run all tests) of the patient for clinical trials at cancer centers in their region.

A prototype to enable physicians, nurses or technicians to enter new protocols has been completed. The system is now in use. It reduces the time required to add a new trial or protocol to approximately 1 hour. It enables non-computer scientists to add trial/protocols to the system. This knowledge acquisition tool has been designed to minimize/eliminate the cases where similar questions acquiring essentially the same information would have to be asked. This feature has the potential to cause slight changes to the wording of inclusion/exclusion criteria. We believe that this change is minor and will have no effect on IRB approval. However, this year we will have new protocols entered using existing questions and go back to the IRB board to discuss any changes in criteria wording to fit existing questions within the system. An example would be a protocol in which there are two questions which ask if a test value is greater than some threshold and then a separate question that asks if it is less than some threshold, versus a single question which asks if a test is in some range. We believe that such a change is trivial, but this must be addressed in practice and we will evaluate whether it causes review board decisions to potentially change.

5.1 So What

The prototype system shows the potential for allowing community physicians, as well as cancer center physicians, to quickly and cost effectively determine for which clinical trials a patient may be eligible. It holds the promise of enabling greater patient accrual for trials by increasing the awareness of each trial for treating physicians throughout a region. In this next year, we will be evaluating how many patients not eligible for clinical trials were actually missed by clinical practitioners vs. excluded for a particular reason (e.g. it was clear they would not agree) or were offered a trial and declined to enter it.

References

- [1] Savvas Nikiforou. Selection of clinical trials: Knowledge representation and acquisition. Master's thesis, University of South Florida, 2002.
- [2] Savvas Nikiforou, Eugene Fink, Lawrence O. Hall, Dmitry B. Goldgof, and Jeffrey P. Krischer. Knowledge acquisition for clinical-trial selection. In *IEEE International Conference on Systems, Man, and Cybernetics*, 2002. To Appear.
- [3] Princeton Kokku, Lawrence O. Hall, Dmitry B. Goldgof, Eugene Fink, and Jeffrey P. Krischer. A cost-effective agent for clinical trial assignment. In *IEEE International Conference on Systems, Man, and Cybernetics*, 2002. To Appear.

Table 1: Current Patient data : Number of patients checked is 57. Number of currently accruing trials is 7.

Clinical Trial Number	Same Matches	New Matches	Missing Data	Patients Checked	Predicted Eligibility
11132	4	1	1	7	5
11931	1	8	0	57	9
11971	3	0	0	56	3
12100	0	2	0	55	2
12101	4	21	0	55	25
12601	0	1	0	50	1
12775	1	4	0	24	5
Total	13	37	1	57	50

A Cost-Effective Agent for Clinical Trial Assignment

Princeton K. Kokku, Lawrence O. Hall, Dmitry B. Goldgof, Eugene Fink, and Jeffrey P. Krischer

kokku@csee.usf.edu, hall@csee.usf.edu, goldgof@csee.usf.edu,
eugene@csee.usf.edu, jpkrisher@moffitt.usf.edu

Computer Science and Engineering, University of South Florida, Tampa, Florida 33620

Abstract— The purpose of a clinical trial is to evaluate a new treatment procedure. When medical researchers conduct a trial, they recruit participants with appropriate medical histories. To select participants, the researchers analyze medical records of the available patients, which has traditionally been a manual procedure. We describe an intelligent agent that helps to select patients for clinical trials. If the available data are insufficient for choosing patients, the agent suggests additional medical tests and finds an ordering of the tests that reduces their total cost.

Keywords— Medical expert systems, automated diagnosis, clinical trials.

I. INTRODUCTION

A *clinical trial* is an experiment with a new treatment procedure. When medical researchers test a new treatment, they recruit patients with appropriate health problems and medical histories. The selection of patients has traditionally been a manual procedure, and recent studies have shown that clinicians can miss up to 60% of the eligible patients [9, 10, 14, 26, 35, 38].

If the available records do not provide enough data, clinicians perform medical tests as part of the selection process. The costs of most tests have declined over the last decade, but the number of tests has significantly increased [33, 36], which is partially due to inappropriate ordering of tests [1, 25]. Clinicians can reduce the cost by first requiring inexpensive tests and then using their results to avoid some expensive tests; however, finding the right ordering may be a complex problem.

The purpose of the described work is to automate the selection of patients for clinical trials and minimize the cost of related tests. We have developed an agent that identifies appropriate trials for each patient, and built a knowledge base for breast-cancer trials.

II. PREVIOUS WORK

Researchers began to work on medical expert systems in the early seventies. Shortliffe *et al.* developed the MYCIN system, which diagnosed bacterial diseases [5, 30, 31]. Its knowledge base consisted of *if-then* rules, which allowed for the analysis of symptoms and evaluation of the certainty of the diagnosis. Experiments showed that MYCIN correctly diagnosed common diseases, which led to the development of other medical systems [5, 19], such as NEOMYCIN, PUFF, CENTAUR, and VM. Shortliffe *et al.* created a system for selecting chemotherapy treatments, called ONCOCIN [32], which also evolved from MYCIN.

Lucas *et al.* constructed a rule-based system for diagnosing liver and biliary-tract diseases [16], but it often

gave an incorrect diagnosis [12, 23]. Korver and Lucas converted the initial system into a Bayesian network, which improved its performance [13, 15].

Musen *et al.* built a rule-based system, called EON, that selected AIDS patients for clinical trials [20]. Ohno-Machado *et al.* developed the AIDS² system, which also assigned AIDS patients to clinical trials [21]. They integrated logical rules with Bayesian networks, which helped to make decisions in the absence of some data.

Bouaud *et al.* created a cancer expert system, called ONCODOC, that suggested alternative clinical trials for each patient and allowed a physician to choose among them [3, 4]. Séroussi *et al.* used ONCODOC to select participants for clinical trials at two hospitals, which helped to increase the number of selected patients by a factor of three [27, 28, 29].

Hammond and Sergot created the OaSiS architecture [11], which combined the techniques from earlier systems, including EON and ONCOCIN. Smith *et al.* built a system that assisted a clinician in selecting medical tests and reducing their number and cost [17, 18, 33].

Fallowfield *et al.* studied how physicians selected cancer patients for clinical trials, and compared manual and automatic selection [8]. They showed that expert systems could improve the selection accuracy; however, their study also revealed that physicians were reluctant to use these systems. Carlson *et al.* conducted similar studies with AIDS trials, and also concluded that expert systems could lead to a more accurate selection [6].

Theocharous developed a Bayesian system that selected clinical trials for cancer patients [24, 34]. It learned conditional probabilities of medical-test outcomes and evaluated the probability of a patient's eligibility for each trial. On the negative side, the available medical records were often insufficient for learning accurate probabilities. Furthermore, when adding a new clinical trial, the user had to change the structure of the underlying Bayesian network.

To address these problems, Bhanja *et al.* built a rule-based system for the same task [2]. We have continued that work, extended the system, and added a mechanism for reducing costs involved in patient selection.

III. EXAMPLE

We have developed an intelligent agent that helps to select clinical trials for eligible patients. It prompts a clinician to enter the results of medical tests, and identifies appropriate trials. If the available records do not provide enough data, the agent suggests additional tests.

In Figure 1(a), we give a simplified example of eligibil-

(a) Eligibility criteria

1. The patient is female.
2. She is at most forty-five years old.
3. Her cancer stage is II or III.
4. Her cancer is not invasive.
5. At most three lymph nodes have tumor cells.
6. Either
 - the patient has no cardiac arrhythmias, or
 - all tumors are smaller than 2.5 centimeters.

(b) Tests and questions

General information

What is the patient's sex?
What is the patient's age?

Mammogram, Cost is \$150

What is the cancer stage?
Does the patient have invasive cancer?

Biopsy, Cost is \$300

What is the cancer stage?
How many lymph nodes have tumor cells?
What is the greatest tumor size?

Electrocardiogram, Cost is \$200

Does the patient have cardiac arrhythmias?

Fig. 1. Example of eligibility criteria, tests, and questions.

(a) Acceptance	(b) Rejection
$sex = \text{FEMALE}$ and $age \leq 45$ and $stage \in \{\text{II, III}\}$ and $invasive = \text{NO}$ and $lymph-nodes \leq 3$ and $(arrhythmias = \text{NO} \text{ or } tumor-size \leq 2.5)$	$sex = \text{MALE}$ or $age > 45$ or $cancer \in \{\text{I, IV}\}$ or $invasive = \text{YES}$ or $lymph-nodes > 3$ or $(arrhythmias = \text{YES} \text{ and } tumor-size > 2.5)$

Fig. 2. Logical expressions for the criteria in Figure 1(a).

ity criteria for a clinical trial. This trial is for young and middle-aged women with a noninvasive cancer at stage II or III. When testing a patient's eligibility, a clinician has to order three medical tests (Figure 1b).

The agent first prompts a clinician to enter the patient's sex and age. If the patient satisfies the corresponding conditions, the agent asks for the mammogram results and verifies Conditions 3 and 4; then, it requests the biopsy and electrocardiogram data. If the patient's records already include some test results, the clinician can answer the corresponding questions while entering the personal data, before the agent selects test procedures. For example, if the records indicate that the cancer stage is IV, the clinician can enter the stage along with sex and age, and then the agent immediately determines that the patient is ineligible for this trial.

IV. KNOWLEDGE BASE

The agent's knowledge base includes questions, medical tests, and logical expressions that represent eligibility criteria for each trial. We give a simplified example of tests and questions in Figure 1(b), and logical expressions in Figure 2.

$$\left(\begin{array}{l} sex = \text{FEMALE} \text{ and} \\ age \leq 45 \text{ and} \\ stage \in \{\text{II, III}\} \text{ and} \\ invasive = \text{NO} \text{ and} \\ lymph-nodes \leq 3 \text{ and} \\ arrhythmias = \text{NO} \end{array} \right) \text{ or } \left(\begin{array}{l} sex = \text{FEMALE} \text{ and} \\ age \leq 45 \text{ and} \\ stage \in \{\text{II, III}\} \text{ and} \\ invasive = \text{NO} \text{ and} \\ lymph-nodes \leq 3 \text{ and} \\ tumor-size \leq 2.5 \end{array} \right)$$

Fig. 3. Disjunctive normal form of the acceptance expression.

The agent supports three types of questions; the first type takes a yes/no response, the second is multiple choice, and the third requires a numeric answer. For example, the cancer stage is a multiple-choice question, and the tumor size is a numeric question. The description of a medical test includes the test name, dollar cost, and list of questions that can be answered based on the test results (Figure 1).

We encode the eligibility for a clinical trial by a logical expression that does not have negations, called the *acceptance expression*. It includes variables that represent medical data, as well as equalities, inequalities, "set-element" relations, conjunctions, and disjunctions (Figure 2a). In addition, the agent uses the logical complement of the eligibility criteria, called the *rejection expression*, which also does not have negations (Figure 2b). It describes the conditions that make a patient ineligible for the trial.

The agent collects data until it can determine which of the two expressions is TRUE. For instance, if a patient's sex is MALE, then the rejection expression in Figure 2(b) is TRUE, and the agent immediately determines that this trial is inappropriate. If the sex is FEMALE, the agent asks more questions.

If the knowledge base includes multiple clinical trials, the agent checks a patient's eligibility for each of them. It first asks for the tests related to multiple trials, and then requests additional tests for specific trials. After getting each new answer, the agent re-evaluates the patient's eligibility for each trial.

V. ORDER OF TESTS

If a patient's records do not include enough data, the agent asks for additional tests; for example, if the records do not provide data for the eligibility criteria in Figure 1, the agent asks for the mammogram, biopsy, and electrocardiogram. The total cost of tests may depend on their order; for instance, if we begin with the mammogram, and it shows that the cancer stage is IV, then we can immediately reject the trial in Figure 1 and avoid the more expensive tests.

We have explored heuristics for ordering the tests, based on the test costs and the structure of acceptance and rejection expressions. The heuristics use a disjunctive normal form of these expressions; that is, each expression must be a disjunction of conjunctions. For example, the rejection expression in Figure 2(b) is in disjunctive normal form, whereas the acceptance expression in Figure 2(a) is not. If the system uses ordering heuristics, it converts this acceptance expression into the disjunctive normal form shown in Figure 3.

The agent chooses the order of tests that reduces their expected cost. After getting the results of the first test, it re-evaluates the need for the other tests and revises their ordering. The choice of the first test is based on three criteria. The agent scores all required tests according to these criteria, computes a linear combination of the three scores for every test, and chooses the test with the highest score.

1. *Cost of the test.* The agent prefers cheaper tests. For instance, it may start with the mammogram, which is cheaper than the other two tests in Figure 1.

2. *Number of clinical trials that require the test.* When the agent checks a patient's eligibility for several trials, it prefers tests that provide data for the largest number of trials. For example, if the electrocardiogram gives data for two different trials, the agent may prefer it to the mammogram despite its higher cost.

3. *Number of clauses that include the test results.* The agent prefers the tests that provide data for the largest number of clauses in the acceptance and rejection expressions. For example, the mammogram data affect both clauses of the acceptance expression in Figure 3 and two clauses of the rejection expression in Figure 1(b). On the other hand, the electrocardiogram affects only one clause of the acceptance expression and one clause of the rejection expression; thus, the agent should order it after the mammogram.

VI. USER INTERFACE

The agent includes a web-based interface that allows clinicians to enter patients' data through remote computers; the interface consists of five screens (Figure 4).

The start screen is for adding and retrieving patients (Figure 5). After a user enters a patient's name, the agent displays a list of the available trials (Figure 6). The user can choose a subset of these trials, and then the agent checks eligibility only for the selected trials. The next screen is for basic personal and medical data, such as sex, age, and cancer stage (Figure 7).

After the agent gets the basic data, it prompts the user for medical information related to specific trials (Figure 8). When the user enters medical data, the agent continuously re-evaluates the patient's eligibility and shows the decision for each trial. If the patient is ineligible for some trials, the user can find out the reasons by clicking the "Why" button. The interface also includes a screen for the review and modification of the previous answers, similar to the screen in Figure 8.

VII. EXPERIMENTS

We have built a knowledge base for the breast-cancer clinical trials at the H. Lee Moffitt Cancer Center, applied the agent to retrospective data from 187 past patients and 57 current patients, and compared the results with manual selection by clinicians at the cancer center.

We summarize the results for the past patients in Table I, and the results for the current patients in Table II. The "same matches" column includes the number of patients who have been selected by both human clinicians and the automated agent. The "new matches" column gives the number of patients who have been matched

TABLE I
RESULTS OF MATCHING 187 PAST PATIENTS.

Clinical Trial	Same Matches	New Matches	Missing Data
10822	10	5	0
10840	0	19	3
11072	48	26	19
11378	4	19	3
11992	5	6	0
12100	8	20	13
12101	20	30	0

TABLE II
RESULTS OF MATCHING 57 CURRENT PATIENTS.

Clinical Trial	Same Matches	New Matches	Missing Data
11132	4	1	1
11971	3	0	0
12100	0	2	0
12101	4	21	0
12601	0	1	0
11931	1	8	0
12775	1	4	0

by the agent but potentially missed by human clinicians. Finally, the last column shows the number of patients whose available records are incomplete. Clinicians have found trials for these patients, but the agent cannot identify these matches because of missing data. The agent has found a number of matches potentially missed by human clinicians; thus, it can help to recruit more patients for clinical trials.

In Table III, we give the mean test costs with and without the ordering heuristics for the 187 past patients. The results show that the implemented heuristics reduce the costs by more than a factor of two.

VIII. SCALABILITY

The time complexity of evaluating the acceptance and rejection expressions is linear in their size. Experiments on a Sun Ultra 10 have shown that the evaluation takes about 0.02 seconds per question, and the time is linear in the number of questions. Typical eligibility conditions for a clinical trial include ten to thirty questions; thus, the evaluation time is 0.2 to 0.6 seconds per trial.

TABLE III
COST SAVINGS BY TEST REORDERING.

Clinical Trial	Average Dollar Cost	
	Without Test Reordering	With Test Reordering
10822	\$20	\$8
10840	\$0	\$0
11072	\$556	\$194
11378	\$34	\$0
11992	\$87	\$34
12100	\$0	\$0
12101	\$24	\$22

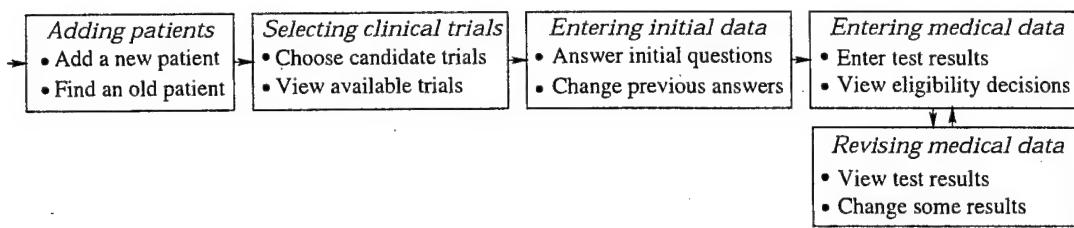


Fig. 4. Entering a patient's data. The web-based interface for data entry consists of five screens. We show these screens by rectangles and the transitions between them by arrows.

Fig. 5. Adding new patients and retrieving existing patients.

Fig. 6. Selecting clinical trials.

Fig. 7. Entering basic information for a patient.

PROTOCOL	STATUS	QUESTIONS REMAINING	PERCENTAGE OF QUESTIONS ANSWERED
001	More Information Needed	14	17 <input checked="" type="checkbox"/> Why?
001	More Information Needed	14	17
002	Eligible	23	14
003	Ineligible	27	10
Does the patient have invasive cancer?		Does the patient have cardiac arrhythmias?	
<input type="radio"/> Yes <input checked="" type="radio"/> No <input type="radio"/> Defer		<input type="radio"/> Yes <input checked="" type="radio"/> No <input type="radio"/> Defer	
Does the patient have recurrent cancer?		Does the patient have congenital heart disease?	
<input type="radio"/> Yes <input checked="" type="radio"/> No <input type="radio"/> Defer		<input type="radio"/> Yes <input checked="" type="radio"/> No <input type="radio"/> Defer	
PROCESS <input type="button" value="Click to submit your answers"/>		REVIEW <input type="button" value="Click to review and change your answers"/>	

Fig. 8. Entering medical data.

(a) Eligibility criteria

1. The patient is female.
2. She is at most forty-five years old.
3. Either
 - her cancer is not invasive, or
 - her cancer is not recurrent.
4. Either
 - at most three lymph nodes have tumor cells, or
 - all tumors are smaller than 2.5 centimeters.
5. Either
 - the patient has no cardiac arrhythmias, or
 - the patient has no congenital heart disease.

(b) Acceptance expression

sex = FEMALE and
age \leq 45 and
(*invasive* = NO or *recurrent* = NO) and
(*lymph-nodes* \leq 3 or *tumor-size* \leq 2.5) and
(*arrhythmias* = NO or *congenital* = NO)

(c) Reduced expression

sex = FEMALE and
age \leq 45 and
invasive-and-recurrent = NO and
(*lymph-nodes* \leq 3 or *tumor-size* \leq 2.5) and
arrhythmias-and-congenital = NO

Fig. 9. Reducing the number of disjunctions. The conversion of the eligibility criteria (a) into a logical expression (b) leads to an explosion in the size of the corresponding disjunctive normal form. We can prevent the explosion by replacing some disjunctions with single questions (c).

The linear scalability is an important advantage over Bayesian systems, which do not scale to a large number of clinical trials [7, 21, 23]. The authors of these systems have reported that the sizes of the underlying networks are superlinear in the number of trials [22, 37], and the training time is superlinear in the network size [24, 34].

If the agent uses the cost-reduction heuristics, it converts the acceptance and rejection expressions into disjunctive normal form, which can potentially lead to an explosion in their size. For example, if eligibility conditions are as shown in Figure 9(a), the agent initially generates the expression in Figure 9(b). If the agent converts it to disjunctive normal form, the resulting expression consists of eight clauses.

Although the conversion may result in impractically large expressions, experiments have shown that this problem does not arise in practice because the number of nested disjunctions is usually small. Furthermore, we can eliminate some disjunctions by combining their elements into longer questions. For instance, we can represent Condition 3 in Figure 9(a) by a single question: "Does the patient have both invasive and recurrent cancer?" If we apply this modification to Conditions 3 and 5, then we obtain the expression in Figure 9(c), and its conversion to disjunctive normal form results in an expression with two clauses.

IX. CONCLUDING REMARKS

We have developed an agent that automatically assigns patients to clinical trials. We have described the representation of selection criteria, heuristics for ordering of tests, and a web-based interface for entering patients' data, which will enable physicians across the country to access a central repository of clinical trials.

Experiments have confirmed that the agent has the potential to find more participants for clinical trials. They have also shown that the ordering of medical tests affects their overall cost, and the implemented heuristics can reduce the cost of finding trial participants. The heuristics do not account for the probabilities of possible test results, and we plan to add probabilistic reasoning as part of the future work.

Acknowledgments: This work has been partially supported by the Breast Cancer Research Program of the U.S. Army Medical Research and Materiel Command under contract DAMD17-00-1-0244, and by the H. Lee Moffitt Cancer Center.

REFERENCES

- [1] D. Bareford and A. Hayling. Inappropriate use of laboratory services: Long term combined approach to modify request patterns. *British Medical Journal*, 301(6764):1305-1307, 1990.
- [2] Sanjukta Bhanja, Lynn M. Fletcher, Lawrence O. Hall, Dmitry B. Goldgof, and Jeffrey P. Krischer. A qualitative expert system for clinical trial assignment. In *Proceedings of the Eleventh International Florida Artificial Intelligence Research Society Conference*, pages 84-88, 1998.
- [3] Jacques Bouaud, Brigitte Séroussi, Éric-Charles Antoine, Mary Gozy, David Khayat, and Jean-François Boisvieux. Hypertextual navigation operationalizing generic clinical practice guidelines for patient-specific therapeutic decisions. *Journal of the American Medical Informatics Association*, 5(suppl.):488-492, 1998.
- [4] Jacques Bouaud, Brigitte Séroussi, Éric-Charles Antoine, Laurent Zelek, and Marc Spielmann. Reusing ONCODOC, a guideline-based decision support system, across institutions: A successful experiment in sharing medical knowledge. In *Proceedings of the American Medical Informatics Association Annual Symposium*, volume 7, 2000.
- [5] Bruce G. Buchanan and Edward H. Shortliffe. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, Reading, MS, 1984.
- [6] Robert W. Carlson, Samson W. Tu, Nancy M. Lane, Tze L. Lai, Carol A. Kemper, Mark A. Musen, and Edward H. Shortliffe. Computer-based screening of patients with HIV/AIDS for clinical trial eligibility. *Online Journal of Current Clinical Trials*, 4(179), 1995.
- [7] Francisco J. Díez, José Mira, E. Iturralde, and S. Zubillaga. DIAVAL, a Bayesian expert system for echocardiography. *Artificial Intelligence in Medicine*, 10(1):59-73, 1997.
- [8] Lesley Fallowfield, D. Ratcliffe, and Robert Souhami. Clinicians' attitudes to clinical trials of cancer therapy. *European Journal of Cancer*, 33(13):2221-2229, 1997.
- [9] John H. Gennari and Madhu Reddy. Participatory design and an eligibility screening tool. In *Proceedings of the American Medical Informatics Association Annual Fall Symposium*, pages 290-294, 2000.
- [10] Carolyn Cook Gotay. Accrual to cancer clinical trials:

Directions from the research literature. *Social Science and Medicine*, 33(5):569–577, 1991.

[11] Peter Hammond and Marek J. Sergot. Computer support for protocol-based treatment of cancer. *Journal of Logic Programming*, 26(2):93–111, 1996.

[12] M. Korver and A. R. Janssens. Development and validation of HEPAR, an expert system for the diagnosis of disorders of the liver and biliary tract. *Medical Informatics*, 16(3):259–270, 1993.

[13] M. Korver and Peter J. F. Lucas. Converting a rule-based expert system into a belief network. *Medical Informatics*, 18(3):219–241, 1993.

[14] Cyrus Kotwall, Leo J. Mahoney, Robert E. Myers, and Linda Decoste. Reasons for non-entry in randomized clinical trials for breast cancer: A single institutional study. *Journal of Surgical Oncology*, 50:125–129, 1992.

[15] Peter J. F. Lucas. Refinement of the HEPAR expert system: Tools and techniques. *Journal of Artificial Intelligence in Medicine*, 6(2):175–188, 1994.

[16] Peter J. F. Lucas, R. W. Segar, and A. R. Janssens. HEPAR: An expert system for the diagnosis of disorders of the liver and the biliary tract. *Liver*, 9:266–275, 1989.

[17] Michael D. McNeely and Beverly J. Smith. An interactive expert system for the ordering and interpretation of laboratory tests to enhance diagnosis and control utilization. *Canadian Medical Informatics*, 2(3):16–19, 1995.

[18] Ian R. Morrison, B. A. Schaefer, and Beverly J. Smith. Knowledge acquisition: The ACQUIRE approach. In *Proceedings of the First Semi-Annual Conference in Policy Making and Knowledge Systems*, 1991.

[19] Mark A. Musen. *Automated Generation of Model-Based Knowledge Acquisition Tools*. Morgan Kaufmann, San Mateo, CA, 1989.

[20] Mark A. Musen, Samson W. Tu, Amar K. Das, and Yuval Shahar. EON: A component-based approach to automation of protocol-directed therapy. *Journal of the American Medical Informatics Association*, 3(6):367–388, 1996.

[21] Lucila Ohno-Machado, Eduardo Parra, Suzanne B. Henry, Samson W. Tu, and Mark A. Musen. AIDS²: A decision-support tool for decreasing physicians' uncertainty regarding patient eligibility for HIV treatment protocols. In *Proceedings of the Seventeenth Annual Symposium on Computer Applications in Medical Care*, pages 429–433, 1993.

[22] Agnieszka Onisko, Marek J. Druzdzel, and Hanna Wasyluk. Learning Bayesian network parameters from small data sets: Application of noisy-OR gates. In *Proceedings of the Workshop on Bayesian and Causal Networks: From Inference to Data Mining*, 2000.

[23] Agnieszka Onisko, Mark J. Druzdzel, and Hanna Wasyluk. Application of Bayesian belief networks to diagnosis of liver disorders. In *Proceedings of the Third Conference on Neural Networks and Their Applications*, pages 730–736, 1997.

[24] Constantinos Papaconstantinou, Georgios Theocharous, and Sridhar Mahadevan. An expert system for assigning patients into clinical trials based on Bayesian networks. *Journal of Medical Systems*, 22(3):189–202, 1998.

[25] Franco Ferraro, Paolo Rossi, Carlo Liva, Adolfo Bulfoni, G. Ganzini, and Adriano Giustinelli. Inappropriate emergency test ordering in a general hospital: Preliminary reports. *Quality Assurance Health Care*, 4:77–81, 1992.

[26] Brigitte Séroussi, Jacques Bouaud, and Éric-Charles Antoine. Enhancing clinical practice guideline compliance by involving physicians in the decision process. In Werner Horn, Yuval Shahar, Greger Lindberg, Steen Andreassen, and Jeremy C. Wyatt, editors, *Artificial Intelligence in Medicine*, pages 76–85. Springer-Verlag, Berlin, Germany, 1999.

[27] Brigitte Séroussi, Jacques Bouaud, and Éric-Charles Antoine. Users' evaluation of ONCODOC, a breast cancer therapeutic guideline delivered at the point of care. *Journal of the American Medical Informatics Association*, 6(5):384–389, 1999.

[28] Brigitte Séroussi, Jacques Bouaud, and Éric-Charles Antoine. ONCODOC: A successful experiment of computer-supported guideline development and implementation in the treatment of breast cancer. *Artificial Intelligence in Medicine*, 22(1):43–64, 2001.

[29] Brigitte Séroussi, Jacques Bouaud, Éric-Charles Antoine, Laurent Zelek, and Marc Spielmann. Using ONCODOC as a computer-based eligibility screening system to improve accrual onto breast cancer clinical trials. In Silvana Quaglini, Pedro Barahona, and Steen Andreassen, editors, *Artificial Intelligence in Medicine*, pages 421–430. Springer-Verlag, Berlin, Germany, 2001.

[30] Edward H. Shortliffe. *MYCIN: A Rule-Based Computer Program for Advising Physicians Regarding Antimicrobial Therapy Selection*. PhD thesis, Computer Science Department, Stanford University, 1974.

[31] Edward H. Shortliffe, Randall Davis, Stanton G. Axline, Bruce G. Buchanan, Cordell C. Green, and Stanley Cohen. Computer-based consultations in clinical therapeutics: Explanation and rule acquisition capabilities of the MYCIN system. *Computers and Biomedical Research*, 8:303–320, 1975.

[32] Edward H. Shortliffe, A. Carlisle Scott, Miriam B. Bischoff, William van Melle, and Charlotte D. Jacobs. ONCOCIN: An expert system for oncology protocol management. In *Proceedings of the Seventh International Joint Conference in Artificial Intelligence*, pages 876–881, 1981.

[33] Beverly J. Smith and Michael D. McNeely. The influence of an expert system for test ordering and interpretation on laboratory investigations. *Clinical Chemistry*, 45(8):1168–1175, 1999.

[34] Georgios Theocharous. An expert system for assigning patients into clinical trials based on Bayesian networks. Master's thesis, Computer Science and Engineering Department, University of South Florida, 1996.

[35] Samson W. Tu, Carol A. Kemper, Nancy M. Lane, Robert W. Carlson, and Mark A. Musen. A methodology for determining patients' eligibility for clinical trials. *Journal of Methods of Information in Medicine*, 32(4):317–325, 1993.

[36] Carl Van Walraven and C. David Naylor. Do we know what inappropriate laboratory utilization is? A systematic review of laboratory clinical audits. *Journal of the American Medical Association*, 280(6):550–558, 1998.

[37] Haiqin Wang and Marek J. Druzdzel. User interface tools for navigation in conditional probability tables and elicitation of probabilities in bayesian networks. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, pages 617–625, 2000.

[38] Salim Yusuf, Peter Held, K. K. Teo, and Elizabeth R. Toretsky. Selection of patients for randomized controlled trials: Implications of wide or narrow eligibility criteria. *Statistics in Medicine*, 9:73–86, 1990.

Knowledge Acquisition for Clinical-Trial Selection

Savvas Nikiforou, Eugene Fink, Lawrence O. Hall, Dmitry B. Goldgof, and Jeffrey P. Krischer

nikiforo@csee.usf.edu, eugene@csee.usf.edu, hall@csee.usf.edu,
goldgof@csee.usf.edu, jpkrischer@moffitt.usf.edu

Computer Science and Engineering, University of South Florida, Tampa, Florida 33620

Abstract— When medical researchers test a new treatment procedure, they recruit patients with appropriate medical histories. An experiment with a new procedure is called a clinical trial. The selection of patients for clinical trials has traditionally been a labor-intensive task, which involves the matching of medical records with a list of eligibility criteria, and studies have shown that clinicians can miss up to 60% of the eligible patients. A recent project at the University of South Florida has been aimed at the automation of this task. We have developed an intelligent agent that selects trials for eligible patients. We report the work on the representation and entry of the related knowledge about clinical trials. We describe the structure of the agent's knowledge base and the interface for adding new trials.

Keywords— Knowledge representation, medical expert systems, user interfaces.

I. INTRODUCTION

Cancer causes 550,000 deaths in the United States every year, and the treatment of cancer is an active research area. Medical experts explore new treatment methods, such as drugs, surgery techniques, and radiation therapies. An experiment with a new treatment procedure is called a *clinical trial*. When researchers conduct a trial, they recruit patients with an appropriate cancer type and medical history. The selection of patients has traditionally been a manual procedure, and studies have shown that clinicians can miss up to 60% of the eligible patients [12, 22, 30].

A recent project at the University of South Florida has been aimed at automatic selection of patients for clinical trials. We have developed an intelligent agent that prompts a clinician for a patient's data and identifies all matching trials [1, 11]. It includes a knowledge base with information about available clinical trials, criteria for selecting patients, and related medical tests.

We report the work on a web-based interface that enables a clinician to enter new trials without the help of a programmer. We have used the interface to build a knowledge base for clinical trials at the Moffitt Cancer Center, located at the University of South Florida. We review the previous work on medical expert systems (Section II), explain the knowledge representation in the developed agent (Section III), and describe the interface for adding new knowledge (Section IV).

II. PREVIOUS WORK

Researchers began to work on medical applications of artificial intelligence in the early seventies. Shortliffe and his colleagues developed the MYCIN system,

which diagnosed bacterial diseases [5, 25, 26]. Experiments showed the effectiveness of MYCIN, which led to the development of other medical systems [5, 14], such as NEOMYCIN, PUFF, CENTAUR, and VM.

Musen *et al.* built a rule-based system, called EON, that selected AIDS patients for clinical trials [17]. Ohno-Machado *et al.* developed the AIDS² system, which also assigned AIDS patients to clinical trials [19]. Bouaud *et al.* created a cancer expert system, called ONCODOC, that suggested alternative trials for each patient and allowed a physician to choose among them [3, 4]. Séroussi used ONCODOC to select participants for clinical trials at two hospitals, which helped to increase the number of selected patients by a factor of three [23, 24].

Early expert systems did not have knowledge-acquisition tools, and programmers hand-coded the related rules. To simplify knowledge entry, researchers implemented specialized tools for some systems [13, 15].

Eriksson pointed out the need for tools that would allow efficient knowledge acquisition, and described a system for building such tools [6]. Tallis *et al.* developed a library of scripts for modifying knowledge bases, which helped to enforce the consistency of the modified knowledge [7, 27, 28, 29]. Kim and Gil considered the use of scripts for building new knowledge-acquisition tools, and created a system for evaluating these tools [9, 10]. Blythe *et al.* designed a general knowledge-acquisition interface based on previous techniques [2].

Musen developed the PROTÉGÉ environment for creating knowledge-acquisition tools [14, 16], which proved effective for the development of knowledge systems, including the AIDS expert systems [20], asthma treatment selection [8], and elevator-design rules [21].

III. KNOWLEDGE BASE

Physicians at the Moffitt Cancer Center have about 150 clinical trials available for cancer patients. They have identified criteria that determine a patient's eligibility for each trial, and they use these criteria to select trials for eligible patients. Traditionally, physicians have selected trials by a manual analysis of patients' data. The review of resulting selections has shown that they usually do not check all clinical trials and occasionally miss an appropriate trial.

To address this problem, we have built an intelligent agent that helps to select trials for each patient. It prompts a clinician to enter the results of medical tests, and uses them to identify appropriate trials.

In Figure 1(a), we give a simplified example of eligibility criteria for a clinical trial. This trial is for young and

(a) Eligibility criteria

1. The patient is female.
2. She is at most forty-five years old.
3. Her cancer stage is II or III.
4. Her cancer is not invasive.
5. At most three lymph nodes have tumor cells.
6. Either
 - the patient has no cardiac arrhythmias, or
 - all tumors are smaller than 2.5 centimeters.

(b) Tests and questions

General information

What is the patient's sex?
What is the patient's age?

Mammogram, Cost is \$150

What is the cancer stage?
Does the patient have invasive cancer?

Biopsy, Cost is \$300

What is the cancer stage?
How many lymph nodes have tumor cells?
What is the greatest tumor diameter?

Electrocardiogram, Cost is \$200

Does the patient have cardiac arrhythmias?

(c) Eligibility expression.

sex = FEMALE and
age \leq 45 and
cancer-stage \in {II, III} and
invasive-cancer = NO and
lymph-nodes \leq 3 and
(*arrhythmias* = NO or
tumor-diameter \leq 2.5)

Fig. 1. Example of eligibility criteria, tests and questions.

middle-aged women with a noninvasive cancer at stage II or III. When testing a patient's eligibility, a clinician has to order three medical tests (Figure 1b). The agent first prompts the clinician to enter the patient's sex and age. If the patient satisfies the corresponding conditions, the agent asks for the mammogram results and verifies Conditions 3 and 4; then, it requests the biopsy and electrocardiogram data.

The agent's knowledge base includes questions, tests, and logical expressions that represent eligibility for each trial. We give an example of tests and questions in Figure 1(b), and a logical expression in Figure 1(c).

The agent supports three types of questions; the first type takes a yes/no response, the second is multiple choice, and the third requires a numeric answer. For example, the cancer stage is a multiple-choice question, and the tumor diameter is a numeric question. The description of a medical test includes the test name, dollar cost, and list of questions that can be answered based on the test results. For instance, the mammogram in Figure 1 has a cost of \$150, and it allows the answering of two questions. Different tests may answer the same question; for example, both mammogram and biopsy show the cancer stage.

We encode the eligibility for a clinical trial by a logical expression, which may include variables that represent the available medical data, as well as equalities, inequalities, "set-element" relations, conjunctions, and disjunctions. For example, we encode the criteria in Figure 1(a) by the expression in Figure 1(c).

The agent collects data until it can determine whether the eligibility expression is TRUE or FALSE. For instance, if a patient's sex is MALE, then the expression in Figure 1(c) is FALSE, and the agent immediately rejects this trial. If the sex is FEMALE, the agent has to ask more questions. If the knowledge base includes many clinical trials, the agent checks a patient's eligibility for each of them. It first asks for the tests related to multiple trials, and then requests additional tests for specific trials.

IV. ENTERING ELIGIBILITY CRITERIA

We have designed a web-based interface for adding new clinical trials [18], which consists of two main parts; the first part is for adding information about medical tests (Figure 2), and the second is for eligibility criteria (Figure 3). The interface includes ten screens; two of them are "start screens," which can be reached from any other screen. We give an example of entering eligibility criteria, describe the two parts of the interface, and present experiments on its effectiveness.

Example: Suppose that a user needs to enter the criteria shown in Figure 1. First, she utilizes the "Adding tests" screen to enter the three tests (Figure 4). Then, she adds the related questions; to enter questions for a specific test, she selects the test and clicks "Modify" (Figure 4), and the agent displays the "Modifying a test" screen (Figure 5). To add a question, she clicks the appropriate button at the bottom (Figure 5) and then types the question (Figure 6).

After adding the questions for all tests, the user goes to the "Adding clinical trials" screen and initializes a new trial (Figure 7). She gets the "Selecting tests" screen and chooses the tests related to the current trial (Figure 8). Then, she marks relevant questions and the answers that make a patient eligible (Figure 9). If the eligibility criteria include disjunctions, she has to use the screen for composing logical expressions (Figure 10).

Tests and questions: The interface for adding tests and questions includes six screens (Figure 2). The start screen is for viewing the available tests and defining new ones, whereas the other screens are for modifying tests and adding questions.

We show the start screen in Figure 4; its left-hand side allows viewing questions and going to a modification screen. If the user selects a test and clicks "View," the agent shows the questions related to this test. If the user clicks "Modify," it displays the "Modifying a test" screen (Figure 5). The right-hand side of the start screen allows adding a new test by specifying its name and cost.

The "Modifying a test" screen shows the information about a specific test, which includes the test name, cost, and related questions. The user can change the test name and cost; the four bottom buttons allow moving to the screens for adding and deleting questions.

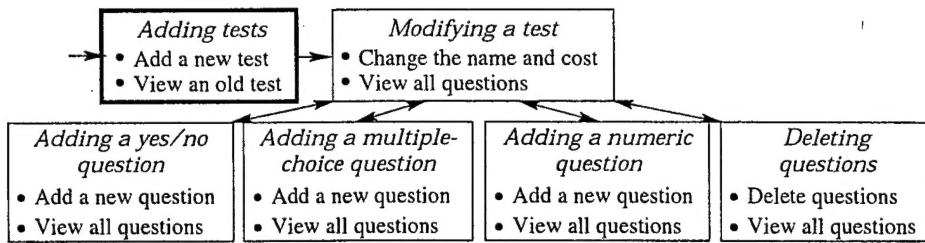


Fig. 2. Entering tests and questions. We show the screens by rectangles and the transitions between them by arrows. The bold rectangle is the start screen.

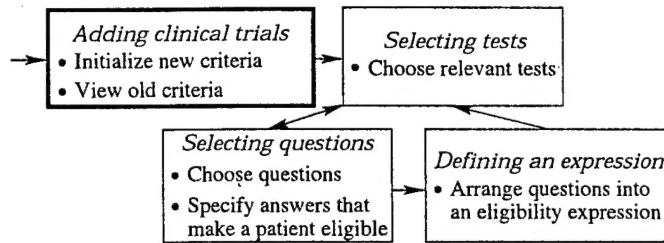


Fig. 3. Entering eligibility criteria.

Current Tests	Add New Test
Mammogram Biopsy	Test Name: <input type="text" value="Electrocardiogram"/> Cost(\$): <input type="text" value="300"/>
<input type="button" value="Modify"/>	<input type="button" value="Add Test"/>

Fig. 4. Adding a new test.

Name: <input type="text" value="Mammogram"/>	Cost(\$): <input type="text" value="150"/>	<input type="button" value="Change"/>	<input type="button" value="Reset"/>
<input type="button" value="Yes/No Question"/> <input type="button" value="Multiple Choice Question"/> <input type="button" value="Numeric Question"/> <input type="button" value="Delete Questions"/>			

Fig. 5. Modifying a test; the bottom buttons are for moving to question-entry screens.

Entering a new Yes/No question

Enter question in the box below

Does the patient have invasive cancer?

Biopsy
 Electrocardiogram

Select other tests that also answer this question

(a) Yes/no question.

Entering a new multiple choice question

Enter question in the box below

What is the cancer stage?

I
 II
 III
 IV

Biopsy
 Electrocardiogram

Select other tests that also answer this question

(b) Multiple-choice question.

Fig. 6. Adding new questions; the user enters a question and answer options.

Protocol Number	Protocol Name
<input type="text" value="001"/>	<input type="text" value="Clinical Trial 1 for breast cancer patients at the Moffitt Cancer Center"/>
<input type="button" value="Add Protocol"/>	<input type="button" value="Clear"/>

Fig. 7. Adding a new clinical trial.

Protocol **?** 001 Clinical Trial 1 for breast cancer patients at the Moffitt Cancer Center

Select Tests **?** Select Questions **?**

General Information	Yes/No Questions
Mammogram	Multiple Choice Questions
Biopsy	Numeric Questions
Electrocardiogram	

Continue **Clear**

Fig. 8. Choosing tests and question types.

Protocol 001 Clinical Trial 1 for breast cancer patients at the Moffitt Cancer Center

Check all | Uncheck all

Yes under "Unknown" means eligible, whereas No means ineligible ?	Unknown
Yes/No Questions	Yes No N/A Yes No
<input checked="" type="checkbox"/> Does the patient have invasive cancer?	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/>
Multiple Choice Questions	
<input checked="" type="checkbox"/> What is the patient's sex?	Options ? Female <input checked="" type="radio"/> Male <input type="radio"/>
<input checked="" type="checkbox"/> What is the cancer stage?	<input type="radio"/> <input type="radio"/> <input type="radio"/>
Numeric Questions	
<input checked="" type="checkbox"/> What is the patient's age?	From: 0 To: 45
<input checked="" type="checkbox"/> How many lymph nodes have tumor cells?	From: 0 To: 3

Simple questions **Combined question** **Clear**

Fig. 9. Selecting questions and answers. The user checks the questions for the current clinical trial and marks the answers that satisfy the eligibility criteria.

1 Does the patient have cardiac arrhythmias?	No
2 What is the greatest tumor diameter?	From 0 To 2.5

Define a logical expression

1 AND OR 2

Continue **Clear**

Update tree

OR

1 Does the patient have cardiac arrhythmias? (No)
2 What is the greatest tumor diameter? (From: 0 To: 2.5)

Fig. 10. Combining questions into a logical expression.

We show the screens for adding yes/no and multiple-choice questions in Figure 6; the screen for numeric questions is similar. The user can enter a new question for the current test, along with a set of allowed answers. If the question is also related to other tests, the user has to mark them in the lower box. The "Deleting questions" screen is for removing old questions.

Eligibility conditions: The mechanism for entering eligibility criteria consists of four screens (Figure 3). The start screen allows the user to initialize a new clinical trial and view the criteria for old trials. If the user needs to modify a clinical trial, the agent first

displays the test-selection screen (Figure 8). The user then chooses related tests and question types, and clicks "Continue" to get the question list.

The next screen (Figure 9) allows the user to select specific questions and mark the answers that make a patient eligible. For a multiple-choice question, the user may specify several eligibility options; for example, a patient may be eligible if her cancer stage is II or III. For a numeric question, the user has to specify a range of values; for instance, a patient may be eligible if her age is between 0 and 45 years. If the user clicks "Simple questions," the agent generates a conjunction of the

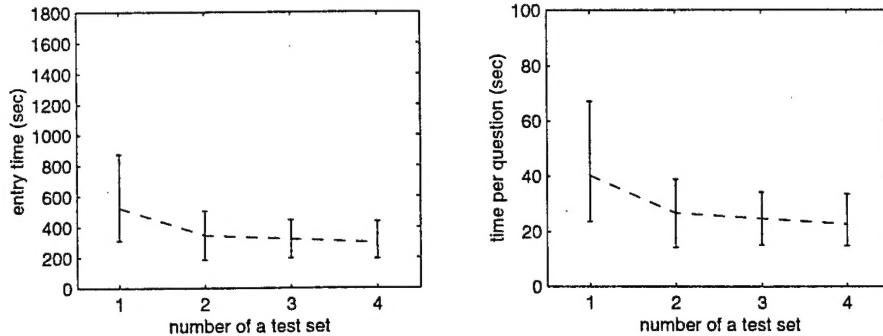


Fig. 11. Entry time for test sets (left) and the mean time per question for each set (right). We plot the average time (dashed lines) and the time of the fastest and slowest users (vertical bars).

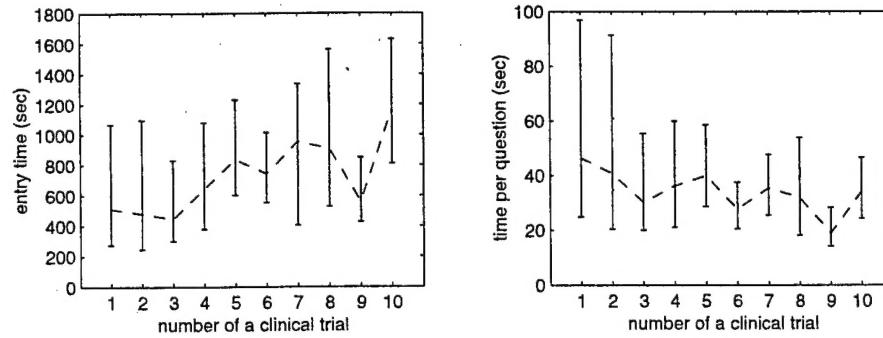


Fig. 12. Entry time for eligibility criteria. We show the average time for each clinical trial and the time per question (dashed lines), along with the performance of the fastest and slowest users (vertical bars).

selected criteria. If the eligibility conditions involve a more complex expression, the user has to click “*Combined question*” and then use the screen for composing logical expressions (Figure 10).

Entry time: We have run experiments with sixteen novice users, who had no prior experience with the interface. First, every user has entered four sets of medical tests; each set has included three tests and ten questions. Then, each user has added eligibility expressions for ten clinical trials used at the Moffitt Cancer Center; the number of questions in an eligibility expression has varied from ten to thirty-five.

We have measured the entry time for each test set and each eligibility expression. In Figure 11, we show the mean time for every test set and the time per question for the same sets. All users have entered the test sets in the same order, from 1 to 4; since they had no prior experience, their performance has improved during the experiment. In Figure 12, we give similar graphs for the entry of eligibility expressions.

The experiments have shown that novices can efficiently use the interface; they quickly learn its full functionality, and their learning curve flattens after about an hour. The average time per question is 31 seconds for the entry of medical tests and 37 seconds for eligibility criteria, which means that a user can enter all 150 cancer trials used at Moffitt in about two weeks.

V. CONCLUDING REMARKS

We have developed knowledge-acquisition tools for an agent that automatically assigns cancer patients to clinical trials. We have described the representation of eligibility criteria and a web-based interface for adding new trials. The experiments have shown that a user can enter a new trial in fifteen to thirty minutes. Novices can use the interface without prior instructions, and they reach their full speed after about an hour. Although cancer research at Moffitt has provided the motivation for this work, the agent is not limited to cancer, and we can use it for trials related to other diseases.

Acknowledgments: This work has been partially supported by the Breast Cancer Research Program of the U.S. Army Medical Research and Materiel Command under contract DAMD17-00-1-0244, and by the Moffitt Cancer Center.

REFERENCES

- [1] Sanjucta Bhanja, Lynn M. Fletcher, Lawrence O. Hall, Dmitry B. Goldgof, and Jeffrey P. Krischer. A qualitative expert system for clinical trial assignment. In *Proceedings of the Eleventh International Florida Artificial Intelligence Research Society Conference*, pages 84–88, 1998.
- [2] Jim Blythe, Jihie Kim, Surya Ramachandran, and Yolanda Gil. An integrated environment for knowledge acquisition. In *Proceedings of the International Conference on Intelligent User Interfaces*, pages 13–20, 2001.

[3] Jacques Bouaud, Brigitte Séroussi, Éric-Charles Antoine, Mary Gozy, David Khayat, and Jean-François Boisvieux. Hypertextual navigation operationalizing generic clinical practice guidelines for patient-specific therapeutic decisions. *Journal of the American Medical Informatics Association*, 5(suppl.):488–492, 1998.

[4] Jacques Bouaud, Brigitte Séroussi, Éric-Charles Antoine, Laurent Zelek, and Marc Spielmann. Reusing ONCODOC, a guideline-based decision support system, across institutions: A successful experiment in sharing medical knowledge. In *Proceedings of the American Medical Informatics Association Annual Symposium*, volume 7, 2000.

[5] Bruce G. Buchanan and Edward H. Shortliffe. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, Reading, MS, 1984.

[6] Henrik Eriksson. Specification and generation of custom-tailored knowledge-acquisition tools. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, volume 1, pages 510–518, 1993.

[7] Yolanda Gil and Marcelo Tallis. A script-based approach to modifying knowledge bases. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pages 377–383, 1997.

[8] Peter D. Johnson and Mark A. Musen. Development of a guideline authoring tool with PROTÉGÉ-II based on the DILEMMA generic protocol and guideline model. Technical Report SMI-96-0620, Stanford Medical Informatics Group, Stanford School of Medicine, 1996.

[9] Jihie Kim and Yolanda Gil. Acquiring problem-solving knowledge from end users: Putting interdependency models to the test. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 223–229, 2000.

[10] Jihie Kim and Yolanda Gil. User studies of an interdependency-based interface for acquiring problem-solving knowledge. In *Proceedings of the International Conference on Intelligent User Interfaces*, pages 165–168, 2000.

[11] Princeton K. Kokku, Lawrence O. Hall, Dmitry B. Goldgof, Eugene Fink, and Jeffrey P. Krischer. A cost-effective agent for clinical trial assignment. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, 2002.

[12] Cyrus Kotwall, Leo J. Mahoney, Robert E. Myers, and Linda Decoste. Reasons for non-entry in randomized clinical trials for breast cancer: A single institutional study. *Journal of Surgical Oncology*, 50:125–129, 1992.

[13] Sandra Marcus and John P. McDermott. SALT: A knowledge acquisition language for propose and revise systems. *Artificial Intelligence*, 39(1):1–37, 1989.

[14] Mark A. Musen. *Automated Generation of Model-Based Knowledge Acquisition Tools*. Morgan Kaufmann, San Mateo, CA, 1989.

[15] Mark A. Musen, David M. Combs, Edward H. Shortliffe, and Lawrence M. Fagan. OPAL: Toward the computer-aided design of oncology advice systems. In Perry L. Miller, editor, *Selected Topics in Medical Artificial Intelligence*, pages 166–180. Springer-Verlag, New York, NY, 1988.

[16] Mark A. Musen, Ray W. Fergerson, William E. Gross, Natalya F. Noy, Monica Crubey, and John H. Gennari. Component-based support for building knowledge-acquisition systems. In *Proceedings of the Conference on Intelligent Information Processing of the International Federation for Information Processing World Computer Congress*, 2000.

[17] Mark A. Musen, Samson W. Tu, Amar K. Das, and Yuval Shahar. EON: A component-based approach to automation of protocol-directed therapy. *Journal of the American Medical Informatics Association*, 3(6):367–388, 1996.

[18] Savvas Nikiforou. Selection of clinical trials: Knowledge representation and acquisition. Master's thesis, Department of Computer Science and Engineering, University of South Florida, 2002.

[19] Lucila Ohno-Machado, Eduardo Parra, Suzanne B. Henry, Samson W. Tu, and Mark A. Musen. AIDS²: A decision-support tool for decreasing physicians' uncertainty regarding patient eligibility for HIV treatment protocols. In *Proceedings of the Seventeenth Annual Symposium on Computer Applications in Medical Care*, pages 429–433, 1993.

[20] Angel R. Puerta, John W. Egger, Samson W. Tu, and Mark A. Musen. A multiple-method knowledge-acquisition shell for the automatic generation of knowledge-acquisition tools. *Knowledge Acquisition*, 4(2):171–196, 1992.

[21] Thomas E. Rothenfluh, John H. Gennari, Henrik Eriksson, Angel R. Puerta, Samson W. Tu, and Mark A. Musen. Reusable ontologies, knowledge-acquisition tools, and performance systems: PROTÉGÉ-II solutions to Sisyphus-2. *International Journal of Human-Computer Studies*, 44(3–4):303–332, 1996.

[22] Brigitte Séroussi, Jacques Bouaud, and Éric-Charles Antoine. Enhancing clinical practice guideline compliance by involving physicians in the decision process. In Werner Horn, Yuval Shahar, Greger Lindberg, Steen Andreassen, and Jeremy C. Wyatt, editors, *Artificial Intelligence in Medicine*, pages 76–85. Springer-Verlag, Berlin, Germany, 1999.

[23] Brigitte Séroussi, Jacques Bouaud, and Éric-Charles Antoine. Users' evaluation of ONCODOC, a breast cancer therapeutic guideline delivered at the point of care. *Journal of the American Medical Informatics Association*, 6(5):384–389, 1999.

[24] Brigitte Séroussi, Jacques Bouaud, and Éric-Charles Antoine. ONCODOC: A successful experiment of computer-supported guideline development and implementation in the treatment of breast cancer. *Artificial Intelligence in Medicine*, 22(1):43–64, 2001.

[25] Edward H. Shortliffe. MYCIN: A Rule-Based Computer Program for Advising Physicians Regarding Antimicrobial Therapy Selection. PhD thesis, Computer Science Department, Stanford University, 1974.

[26] Edward H. Shortliffe, Randall Davis, Stanton G. Axline, Bruce G. Buchanan, Cordell C. Green, and Stanley Cohen. Computer-based consultations in clinical therapeutics: Explanation and rule acquisition capabilities of the MYCIN system. *Computers and Biomedical Research*, 8:303–320, 1975.

[27] Marcelo Tallis. A script-based approach to modifying knowledge-based systems. In *Proceedings of the Tenth Innovative Applications of Artificial Intelligence Conference*, pages 1183–1195, 1998.

[28] Marcelo Tallis and Yolanda Gil. Designing scripts to guide users in modifying knowledge-based systems. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, pages 242–249, 1999.

[29] Marcelo Tallis, Jihie Kim, and Yolanda Gil. User studies of knowledge acquisition tools: Methodology and lessons learned. In *Proceedings of the Twelfth Workshop on Knowledge Acquisition, Modeling and Management*, 1999.

[30] Salim Yusuf, Peter Held, K. K. Teo, and Elizabeth R. Toretsky. Selection of patients for randomized controlled trials: Implications of wide or narrow eligibility criteria. *Statistics in Medicine*, 9:73–86, 1990.